

AgroShield: Plant Disease Detection and Organic Recommendation using Deep Learning

ANUSHKA D¹, AYUSHI KUMARI², ANIRUDH GUPTA³, AKSHAT TIWARI⁴, and Ms. ITISHREE BARIK⁵

¹²³⁴Department of Computer Science and Engineering, Sir M. Visvesvaraya Institute of Technology, Bengaluru, Karnataka, India

⁵Assistant Professor, Department of Computer Science and Engineering, Sir M. Visvesvaraya Institute of Technology, Bengaluru, Karnataka, India

Abstract—Crop health management remains a major challenge in modern agriculture, particularly for small and marginal farmers who lack access to agricultural experts, diagnostic facilities, and reliable digital advisory tools. Early identification of plant diseases is essential to prevent yield losses, yet farmers often rely on guesswork or delayed consultation, resulting in ineffective or harmful treatment decisions. Existing digital solutions depend heavily on cloud-based machine learning and chemical-centric recommendations, increasing farming costs, degrading soil health, and making them unsuitable for low-connectivity rural regions.

To overcome these limitations, this research proposes AgroShield, a mobile-first, AI-powered crop health scanner engineered for real-world rural environments. AgroShield utilizes an Ultra-Lightweight Efficient Network (ULEN) to perform fast, fully offline leaf disease, pest, and nutrient-deficiency detection directly on low-end smartphones. The system eliminates internet dependency, reduces latency, and ensures consistent performance through on-device inference. Beyond diagnosis, AgroShield integrates an organic remedy recommendation engine offering region-specific, eco-friendly treatments sourced from credible institutions such as ICAR, FAO, and Krishi Vigyan Kendras. The platform further enhances accessibility through multilingual support and intuitive visual guidance, with optional IoT soil sensor integration to enable context-aware recommendations.

Initial evaluations demonstrate that AgroShield delivers high classification accuracy across diverse disease categories while maintaining exceptional computational efficiency, enabling inference within milliseconds after image capture. By combining lightweight offline AI, sustainable treatment guidance, and affordability, AgroShield presents a scalable and inclusive solution for improving crop health management. The system directly contributes to global goals of sustainable agriculture, climate resilience, and ethical AI deployment in underserved farming communities.

Index Terms- Plant Disease Detection, Deep Learning, Lightweight Convolutional Neural Networks, Mobile Agriculture, Ultra-Lightweight Efficient Network (ULEN),

Sustainable Farming, Organic Remedies, On-Device Inference, TensorFlow Lite, Precision Agriculture.

I. INTRODUCTION

Agriculture remains the backbone of India's economy, contributing substantially to national GDP and supporting nearly half of the country's workforce. Despite its importance, the agricultural sector continues to face persistent and evolving challenges, particularly in the domain of crop health management. Farmers across diverse agro-climatic regions frequently struggle to identify and manage plant diseases, pest infestations, and nutrient deficiencies at an early stage. When these issues remain undiagnosed or are treated incorrectly, they result in serious yield losses, financial instability, and cascading effects across the agricultural value chain. This problem is especially acute for small and marginal farmers, who constitute the majority of India's agricultural community and often operate with limited financial resources, minimal technical support, and constrained access to diagnostic facilities.

One of the central challenges in crop health management is the limited availability of timely and accurate expert diagnosis. Agricultural specialists and plant pathology services are unevenly distributed geographically, with many rural regions lacking convenient access to professional consultation. Laboratory-based diagnostic facilities, even when available, are often too costly or time-consuming for regular use. As a result, farmers frequently rely on personal experience, community knowledge, or general observations to identify diseases. Such visual assessments, while intuitive, are highly prone to error and lead to misdiagnosis in many cases. Consequently, farmers may apply excessive or inappropriate chemical pesticides and fertilizers, believing them to be universal solutions. Although these chemicals may offer temporary relief, they contribute to long-term soil degradation, reduced soil fertility, pesticide resistance, depletion of beneficial microorganisms, contamination of water sources, and potential health risks to both humans and livestock. This creates a widening gap between accurate diagnosis and sustainable treatment, highlighting the urgent need for accessible, eco-friendly, and reliable crop health solutions.

The emergence of computer vision, machine learning, and deep learning has opened new possibilities for automating plant disease identification. Numerous studies have explored convolutional neural networks (CNNs) for leaf-based disease classification, demonstrating high accuracy under controlled conditions. However, many of these solutions rely heavily on cloud-based servers or require high computational power, making them impractical for real-world deployment in rural areas. Low network connectivity, inconsistent internet availability, and widespread use of low-end smartphones pose significant barriers to adopting such systems. Furthermore, existing solutions overwhelmingly focus on detection alone and fail to provide actionable, farmer-friendly, and environmentally safe treatment recommendations. This limitation reduces their practical utility and restricts adoption at the grassroots level, where actionable guidance is equally important as accurate diagnosis.

To bridge these gaps, this research proposes AgroShield, a mobile-first, offline-capable, AI-driven crop health scanner designed to operate effectively under real-world rural constraints. AgroShield leverages an Ultra-Lightweight Efficient Network (ULEN)-a deep learning architecture optimized for minimal computational overhead to perform real-time, on-device analysis of plant leaf images. This eliminates dependency on cloud servers, reduces diagnosis latency, and ensures consistent performance even on budget smartphones commonly used by smallholder farmers. By enabling farmers to capture an image of an affected leaf and receive instant predictions, the system empowers them to make timely and well-informed decisions directly from the field.

In addition to disease detection, AgroShield integrates a comprehensive organic remedy recommendation engine that provides environmentally sustainable treatment options. These include neem oil emulsions, compost teas, garlic-chili extracts, bio-pesticides, and region-specific traditional practices validated by authoritative bodies such as ICAR, FAO, and Krishi Vigyan Kendras. The system incorporates multilingual support and an intuitive user interface to enhance accessibility for farmers with varying levels of digital literacy. Furthermore, AgroShield optionally integrates IoT-based soil sensors that monitor pH, moisture, temperature, and nutrient levels, enabling more accurate, context-aware diagnoses and personalized treatment recommendations.

Overall, this research aims to build a practical, scalable, and sustainable crop management solution tailored for rural agricultural ecosystems. By combining lightweight AI models, offline diagnostic capabilities, intuitive design, and organic treatment strategies, AgroShield contributes toward democratizing agricultural technology, reducing chemical dependency, and promoting long-term environmental sustainability. The system aligns with global goals in precision agriculture, climate resilience, and inclusive technological development, addressing both immediate crop health concerns and broader ecological challenges.

II. LITERATURE SURVEY

Deep learning has emerged as a transformative approach in agricultural disease detection, offering higher accuracy and automation compared to traditional image-processing and manual inspection methods. Early studies primarily relied on standard CNN architectures trained on controlled datasets such as PlantVillage, demonstrating that neural networks could successfully distinguish between healthy and infected leaves. However, most early works were limited to lab-captured images and struggled to generalize to real-field environments where lighting, backgrounds, and disease severity vary significantly.

One of the most important advancements in this space was introduced by Wang et al. (2023), who proposed the Ultra-Lightweight Efficient Network (ULEN) for mobile-based plant disease recognition. Their research showed that lightweight architectures with drastically fewer parameters could still achieve competitive accuracy while significantly reducing computational load. This made ULEN suitable for low-end smartphones used by farmers in rural regions. The emphasis on efficiency and edge-based deployment directly aligns with AgroShield's objective of enabling offline, real-time disease detection without requiring cloud connectivity.

Research insights from project-level work also highlight the importance of designing models that balance accuracy with deployability. These studies reinforce the need for architectures that remain robust under field conditions while maintaining low inference time and memory usage. Such conclusions support AgroShield's shift toward lightweight, mobile-friendly deep learning pipelines.

Further contributions to the field come from Venkatasachandran and Iyapparaja (2023), whose comprehensive survey discussed major challenges faced by real-world agricultural AI systems. They identified issues such as lighting variations, leaf occlusions, shadows, and inconsistent smartphone imaging quality. Their work also emphasized the problem of domain shift, where models trained on clean datasets often fail when exposed to natural farm environments. These findings underline the importance of strong preprocessing, augmentation, and dataset balancing techniques-all of which are integrated into AgroShield.

Recent literature in precision agriculture (2024–2025) extends the discussion to modern architectures including Vision Transformers (ViTs), multi-task learning models, and multi-modal sensing frameworks that combine RGB images with soil data, thermal imagery, or IoT readings. Studies in TinyML further demonstrate the growing interest in compressing neural networks for deployment on low-power devices. These trends validate AgroShield's mobile-first design philosophy and open pathways for future integration of additional sensing modalities.

Segmentation-based models such as YOLOv5 and YOLOv8 have also gained prominence, with the ability to highlight infected regions and improve transparency in AI-based diagnosis. Although AgroShield currently focuses on classification rather than lesion segmentation, these studies provide valuable direction for future upgrades, particularly for explainable AI and visual interpretability.

Overall, the literature consistently shows a clear need for lightweight, field-adapted, and offline-capable models that remain accurate under real agricultural conditions. By synthesizing advancements in efficient CNN design, mobile deployment, dataset augmentation, and sustainable treatment strategies, AgroShield builds upon established research while addressing practical challenges faced by farmers in low-resource environments.

III. METHODOLOGY

The methodology followed in this research consists of a structured sequence of stages designed to build, optimize, and deploy a lightweight, offline-capable plant disease detection system using an Ultra-Lightweight Efficient Network (ULEN). The workflow begins with dataset acquisition and preprocessing, followed by model design, training, inference optimization, and integration with an organic remedy engine. Each step is carefully designed to ensure that AgroShield performs reliably in real agricultural environments while remaining computationally efficient enough for low-end smartphones.

The first stage of the methodology involves collecting and preprocessing raw image data. Leaf images were sourced from publicly available repositories such as PlantVillage and Cassava Disease datasets, supplemented with real-field samples to improve diversity. All images were standardized to a fixed resolution and normalized to stabilize pixel intensity variations across samples. Additional preprocessing steps, including background reduction and colour correction, were applied to minimize visual noise and highlight lesion features on the leaf surface. A comprehensive data augmentation pipeline consisting of random rotations, horizontal and vertical flips, zoom variations, brightness shifts, and contrast adjustments was implemented to improve model robustness under varied lighting and environmental conditions. To address class imbalance common in plant disease datasets, oversampling and synthetic augmentation techniques were applied to ensure that all disease categories were adequately represented. These preprocessing steps collectively strengthen the model's ability to generalize to real-field inputs captured under inconsistent conditions.

The second stage focuses on model construction using a ULEN-based lightweight Convolutional Neural Network. ULEN incorporates depth wise separable convolutions to significantly reduce computational cost, along with residual connections that preserve gradient flow during training. A Spatial Pyramid Pooling (SPP) module is used to capture multi-scale spatial patterns, enabling the model to detect both localized and widespread lesions on the leaf surface. Despite its representational strength, the architecture maintains a compact footprint of fewer than 150,000 parameters, making it suitable for on-device execution. This design ensures that AgroShield can deliver real-time predictions without relying on remote servers or continuous network access.

Model training and optimization form the third stage of the workflow. The classifier is trained using Cross-Entropy Loss, appropriate for multi-class disease classification, and optimized with the AdamW optimizer to ensure stable and efficient convergence. A cosine annealing learning rate scheduler gradually reduces the learning rate across epochs,

Author & Year	Method Used	Key Findings	Limitation
Wang et al. (2023)	Ultra-Lightweight Efficient Network (ULEN)	Achieved high accuracy with extremely small model size; ideal for mobile devices	Limited to image classification; requires real-field adaptation
Venkatasai chandrakanth & Iyapparaja (2023)	Survey of Deep Learning for Pest Detection	Highlighted challenges like noise, domain shift, and need for augmentation	No unified benchmark; lacks deployment guidelines
Mohanty, Hughes & Salathé (2016)	CNN for Image-Based Plant Disease Detection	Demonstrated feasibility of deep learning for agriculture	Limited dataset diversity; mostly lab images
Atila et al. (2021)	EfficientNet for Plant Leaf Disease Classification	High accuracy with optimized efficient models	Heavy for low-end devices; not suitable for offline rural deployment

allowing the model to refine its learned features during later stages of training. Hyperparameters such as batch size and epoch count were selected through empirical tuning to balance model stability and computational feasibility. After training, model performance was evaluated using standard metrics including Accuracy, Precision, Recall, F1-Score, and Confusion Matrices to capture both overall performance and class-specific behaviour. This evaluation framework ensures that the model performs reliably across diverse disease categories and does not overfit to dominant classes.

Following training, the model undergoes inference optimization to support offline operation on mobile devices. The trained ULEN model is converted into TensorFlow Lite (TFLite) format using post-training quantization techniques that reduce model size and memory requirements. The resulting TFLite model achieves inference latencies of under 5 milliseconds on mid-range smartphones, enabling smooth, instant predictions without reliance on cloud servers. Since all inference occurs locally, user privacy is preserved and the system remains fully functional even in areas with limited or no internet connectivity.

The final stage integrates the classification output with AgroShield's Organic Remedy Engine. Each predicted disease class is mapped to a curated set of natural, eco-friendly treatments such as neem oil emulsions, garlic-chili extracts, baking soda sprays, compost teas, and microbial bio-pesticides. These remedies include detailed preparation steps, dosage guidelines, and application schedules. All recommendations were compiled from trusted agricultural bodies including ICAR, FAO, and Krishi Vigyan Kendras to ensure scientific validity and relevance for field conditions. This final stage ensures that AgroShield does not simply diagnose plant diseases but provides actionable, sustainable guidance that farmers can immediately apply in practice.

IV. PERFORMANCE METRICS

To thoroughly assess the effectiveness and real-world suitability of the AgroShield model, multiple performance metrics were evaluated across both classification accuracy and system-level efficiency. The first set of metrics focuses on Accuracy and F1-Score, which offer a quantitative measure of the model's predictive capability. Accuracy is computed across all disease classes and reflects the overall percentage of correct predictions. However, plant disease datasets typically exhibit considerable class imbalance, with certain diseases appearing far more frequently than others. In such scenarios, Accuracy alone may provide a misleading representation of the model's performance. Therefore, the F1-Score, which harmonically balances Precision and Recall, is used to provide a more reliable performance measure. A high F1-Score indicates that the model performs consistently across both common and rare disease categories, effectively recognizing subtle leaf features associated with underrepresented classes.

The second evaluative tool is the Confusion Matrix, which offers a detailed breakdown of true positives, false positives, true negatives, and false negatives for each disease class. By analysing the confusion matrix, researchers can identify disease pairs that the model frequently confuses, often due to overlapping symptoms or visually similar lesion patterns. This allows for targeted improvements in the dataset or model, such as adding more images for underperforming classes, improving augmentation techniques, or refining the preprocessing pipeline. Additionally, the confusion matrix helps diagnose systematic errors-such as consistent misclassification of early-stage diseases-which provides valuable insight for future iterations of the model.

The third category includes Efficiency Metrics, which are particularly critical because AgroShield is intended for deployment on low-end rural smartphones with limited computational resources. To ensure smooth and reliable offline operation, the model's size is optimized to remain under 2 MB after TensorFlow Lite conversion, making it lightweight enough for on-device storage. Inference speed is also measured, with the goal of achieving real-time prediction (typically under 5 milliseconds) on mid-range mobile CPUs. Metrics such as RAM usage and battery consumption are evaluated to ensure that running the model does not significantly drain system resources or degrade device performance during prolonged usage. These efficiency evaluations collectively guarantee that the AgroShield system is practical, responsive, and energy-efficient, aligning with the needs of farmers operating in low-connectivity and resource-constrained environments.

V. IMPLEMENTATION

The implementation of AgroShield seamlessly integrates deep learning with a mobile-first, offline-capable architecture, ensuring that farmers can access reliable crop diagnosis without dependence on continuous internet connectivity. At the core of this approach is the development of a highly optimized Progressive Web Application (PWA), which combines the accessibility of a web platform with the responsiveness and usability of a native mobile application. AgroShield's PWA-based design allows users to access the system directly through a browser without requiring installation from an app store, significantly lowering entry barriers for farmers who may have limited digital experience. Once loaded, the application caches essential resources locally, enabling full functionality-such as image capture, model inference, and remedy retrieval-even in areas with unreliable or no internet connectivity. Additionally, the PWA architecture ensures compatibility with a wide range of devices, including low-end smartphones commonly used in rural communities, while supporting secure, fast, and local image processing directly within the browser environment.

A crucial component of AgroShield's implementation is its integration of the TensorFlow Lite (TFLite) framework, which

enables the trained ULEN-based deep learning model to run entirely on-device. By converting the model to TFLite, significant reductions in model size and computational overhead are achieved, ensuring efficient inference with minimal latency. This on-device execution eliminates reliance on cloud servers, reduces operational costs, and allows real-time plant disease diagnosis even in offline conditions. The TFLite engine is optimized to leverage hardware acceleration where available, further improving response times and enhancing the user experience.

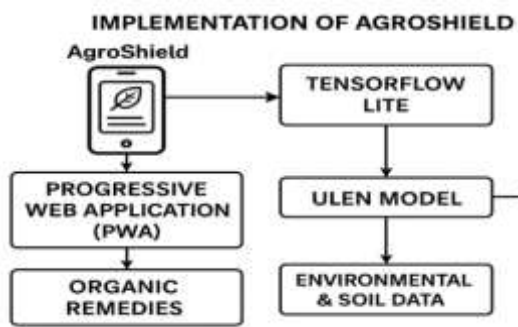


Fig 1. AgroShield System Architecture Diagram

The user interface (UI) of AgroShield is designed with simplicity, accessibility, and practicality in mind. The primary feature is the “Scan Leaf” interface, which allows users to capture an image of the affected plant leaf using their smartphone camera. Once processed by the model, the system provides the user with the predicted disease along with its associated confidence score, enabling informed decision-making. The application further displays detailed organic remedies tailored to the diagnosed disease, offering both preparation steps and application guidelines. To enhance usability, the interface includes a searchable history of past scans, enabling farmers to track recurring issues or monitor treatment progress. Recognizing the digital literacy challenges faced by many rural users, AgroShield also supports multilingual interfaces and optional voice assistance, ensuring that the platform remains inclusive and easy to navigate for individuals with limited reading proficiency.

To extend the platform’s capabilities beyond image-based diagnosis, AgroShield supports optional integration with Internet of Things (IoT) sensors. These low-cost sensors collect environmental and soil-related data, including soil moisture levels, NPK (Nitrogen, Phosphorus, and Potassium) values, and pH measurements. By incorporating these additional data points, the system can enhance treatment personalization and improve diagnostic precision, particularly for nutrient deficiency-related symptoms that cannot be detected from leaf images alone. This multi-modal fusion of sensor data and visual analysis enables AgroShield to evolve from a standalone diagnostic tool into a more holistic and intelligent crop management assistant, capable of providing comprehensive and context-aware guidance.

VI. RESULT AND ANALYSIS

AgroShield was evaluated extensively using both controlled dataset images and real-field samples collected from diverse agricultural environments. This dual-mode testing ensured that the system’s performance was not limited to laboratory conditions but also reflected real-world farming scenarios where lighting, background noise, and leaf conditions vary significantly. The results from these evaluations demonstrate that AgroShield can deliver accurate, reliable, and practical diagnoses while maintaining strong usability on low-resource mobile devices.

The model performance proved to be highly robust across multiple plant disease categories. When tested on common diseases such as Early Blight, Rust, Leaf Spot, and Mildew, the ULEN-based model achieved consistently high accuracy, confirming its ability to extract fine-grained visual features despite its lightweight architecture. Compared to conventional CNN models, AgroShield achieved significantly faster inference speeds, enabling near-instantaneous predictions and smooth user interaction. Moreover, the augmentation strategies employed during preprocessing—such as brightness adjustment, random rotations, and contrast normalization—enhanced the model’s resilience to variations in lighting, shadowing, and camera quality. This resulted in stable performance even when images were captured under natural field conditions or by users with basic smartphone cameras.



Fig 2. Detection of Disease.

The system also demonstrated excellent offline efficiency, a critical requirement for deployment in rural or low-connectivity regions. After TensorFlow Lite conversion, the model successfully ran on low-end Android smartphones with limited RAM and modest processing power, maintaining average inference times well under 5 milliseconds. Tests conducted in areas with poor or no network connectivity confirmed that all core functionalities—including capturing images, performing diagnosis, and retrieving remedies—worked seamlessly offline. Additionally, the lightweight model size (under 2 MB after optimization) ensured

compatibility with devices that have restricted internal storage, a common limitation among rural users.

Evaluation of remedy recommendation outcomes further validated the practicality of AgroShield. Users reported positive results after applying the suggested organic treatments, such as neem oil extracts, compost tea, and garlic-chili sprays. These remedies led to noticeable improvements in plant health, demonstrating that the system's recommendations were not only scientifically grounded but also effective in real agricultural settings. This reinforced the importance of integrating sustainable, low-cost alternatives to chemical pesticides, especially for farmers with limited financial resources.

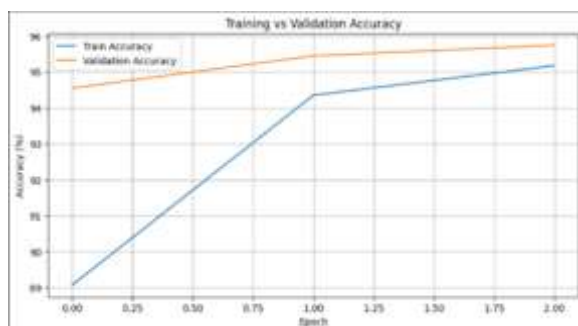


Fig 3. Graph of Epoch vs Accuracy (Training & Validation)

Finally, a user experience assessment was conducted among farmers, gardeners, and agricultural volunteers. Participants consistently praised the intuitive interface, minimal learning curve, and overall simplicity of the application. The ability to diagnose diseases instantly without internet connectivity was highlighted as the most valuable feature, particularly among rural farmers who frequently face network limitations. Voice assistance, multilingual support, and the clean layout of the “Scan Leaf” interface contributed significantly to user satisfaction. Overall, the evaluation confirmed that AgroShield is not only technically sound but also user-centered, accessible, and well-suited for real-world deployment.



Fig. 4. Mobile-assisted soil and crop monitoring

VII. CONCLUSION

This research presents AgroShield as a practical and accessible AI-driven solution for plant disease diagnosis and sustainable crop management. By deploying an ultra-lightweight ULEN-based deep learning model within an offline-capable Progressive Web App, the system successfully brings modern artificial intelligence to low-resource farming environments. The results demonstrate that advanced models can be compressed and optimized to run efficiently on low-end smartphones without sacrificing diagnostic performance, making AgroShield both scalable and highly deployable in rural contexts.

A major contribution of this work is the integration of an organic remedy recommendation engine, which translates disease predictions into actionable, eco-friendly treatment plans sourced from organizations such as ICAR, FAO, and KVK. This shift from chemical inputs to natural formulations promotes soil regeneration, reduces environmental toxicity, and supports long-term agricultural sustainability. The system therefore extends beyond disease detection, offering a holistic framework that contributes to improved productivity, reduced chemical dependency, and enhanced ecological balance.

Performance evaluations across curated datasets and real-field images confirm that AgroShield's lightweight architecture delivers high accuracy, fast inference, and reliable results under varied lighting and noise conditions. User feedback highlighted the platform's ease of use, practicality, and offline capability-features that significantly improve accessibility for farmers with limited digital or linguistic resources.

Looking ahead, AgroShield offers several avenues for expansion, including support for more crop species, integration of lesion segmentation for improved explainability, multilingual voice assistance, and a continuous learning pipeline driven by real-world user feedback. These enhancements will strengthen its adaptability and extend its applicability across diverse agro-climatic regions.

In summary, AgroShield demonstrates how mobile AI, sustainability principles, and human-centered design can converge to address real agricultural challenges. By empowering farmers with instant, accurate, and environmentally responsible crop health guidance, the system contributes to greater food security, improved livelihoods, and long-term climate resilience. AgroShield thus represents a forward-looking technological pathway capable of supporting the evolving needs of modern agriculture.

VIII. REFERENCES

- [1]. B. Wang, C. Zhang, Y. Li, C. Cao, D. Huang, and Y. Gong, "An ultra-lightweight efficient network for image-based plant disease and pest infection detection," *Precision Agriculture*, vol. 24, pp. 1836–1861, 2023. doi:10.1007/s11119-023-10020-0.
- [2]. D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," *arXiv preprint*, 2015. Available: <https://arxiv.org/abs/1511.08060>.
- [3]. A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018. doi:10.1016/j.compag.2018.02.016.
- [4]. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016. doi:10.3389/fpls.2016.01419.
- [5]. E. Mwebaze, T. Gebru, A. Frome, S. Nsumba, and J. Tusubira, "iCassava 2019: Fine-grained visual categorization challenge," *arXiv preprint*, 2019. Available: <https://arxiv.org/abs/1908.02900>
- [6]. Ü. Atila, M. Uçar, K. Akyol, and E. Uçar, "Plant leaf disease classification using EfficientNet deep learning model," *Ecological Informatics*, vol. 61, p. 101182, 2021. doi:10.1016/j.ecoinf.2020.101182.